Notes for "Geometry and topology in many-particle systems' Physics 250, University of California, Berkeley

(Dated: February 10, 2009)

I. MATHEMATICAL PRELIMINARIES

A. An intuitive example of global geometry and topology: Gauss-Bonnet

You may have heard a topologist described as "a mathematican who can't tell the difference between a donut and a coffee cup." As an example of the connections between geometry and topology, we start by discussing an integral that will help us classify two-dimensional compact manifolds (surfaces without boundaries) embedded smoothly in three dimensions. The integral we construct is "topologically invariant" in that if one such surface can be smoothly deformed into another, then the two will have the same value of the integral. The integral can't tell the difference between the surface of a coffee cup and that of a donut, but it can tell that the surface of a donut (a torus) is different from a sphere. Similar connections between global geometry and topology appear frequently in this course.

We start with a bit of local geometry. Given our 2D surface in 3D, we can choose coordinates at any point on the surface so that the $(x, y, z = 0)$ plane is tangent to the surface, which can locally be specified by a single function $z(x, y)$. We choose $(x = 0, y = 0)$ to be the given point, so $z(0, 0) = 0$. The tangency condition is

$$\frac{\partial z}{\partial x}\Big|_{(0,0)} = \frac{\partial z}{\partial y}\Big|_{(0,0)} = 0. \tag{1}$$

Hence we can approximate z locally from its second derivatives:

$$z(x, y) \approx \frac{1}{2}\begin{pmatrix} x & y \end{pmatrix}\begin{pmatrix} \frac{\partial z}{\partial^2 x} & \frac{\partial z}{\partial x \partial y} \\ \frac{\partial z}{\partial y \partial x} & \frac{\partial z}{\partial^2 y} \end{pmatrix}\begin{pmatrix} x \\ y \end{pmatrix} \tag{2}$$

The "Hessian matrix" that appears in the above is real and symmetric. It can be diagonalized and has two real eigenvalues λ_1, λ_2, corresponding to two orthogonal eigendirections in the (x, y) plane. The geometric interpretation of these eigenvalues is simple: their magnitude is an inverse radius of curvature, and their sign tells whether the surface is curving toward or away from the positive z direction in our coordinate system. To see why the first is true, suppose that we carried out the same process for a circle of radius r tangent to the x-axis at the origin. Parametrize the circle by an angle θ that is 0 at the origin and traces the circle counter-clockwise, i.e.,

$$x = r \sin\theta, \quad y = r(1 - \cos(\theta)). \tag{3}$$

Near the origin, we have

$$y = r(1 - \cos(\sin^{-1}(x/r)) = r - (1 - \frac{x^2}{2r^2}) = \frac{x^2}{2r}, \tag{4}$$

which corresponds to an eigenvalue $\lambda = 1/r$ of the matrix in Eq. 2.

Going back to the Hessian, its determinant (the product of its eigenvalues $\lambda_1\lambda_2$) is called the Gaussian curvature and has a remarkable geometric significance. First, consider a sphere of radius r, which at every point has $\lambda_1 = \lambda_2 = 1/r$. Then we can integrate the Gaussian curvature over the sphere's surface,

$$\int_{S^2} \lambda_1\lambda_2 \, dA = \frac{4\pi r^2}{r^2} = 4\pi. \tag{5}$$

Beyond simply being independent of radius, this integral actually gives the same value for any compact manifold that can be smoothly deformed to a sphere.

However, we can easily find a compact manifold with a different value for the integral. Consider the torus made by revolving the circle in Eq. 3, with $r = 1$, around the axis of symmetry $x = t, y = -1, z = 0$, with $-\infty < t < \infty$. To compute the Gaussian curvature at each point, we sketch the calculation of the eigenvalues of the Hessian as follows. One eigenvalue is around the smaller circle, with radius of curvature r: $\lambda_1 = 1/r = 1$. Then the second eigenvalue must correspond to the perpendicular direction, which has a radius of curvature that depends on the angle θ around

the smaller circle (we keep $\theta = 0$ to indicate the point closest to the axis of symmetry). The distance from the axis of symmetry is $2 - \cos\theta$, so we might have guessed $\lambda_2 = (2 - \cos\theta)^{-1}$, but there is an additional factor of $\cos\theta$ that appears because of the difference in direction between the surface normal and this curvature. So our guess is that

$$\lambda_2 = -\frac{\cos\theta}{2 - \cos\theta} \tag{6}$$

As a check and to understand the sign, note that this predicts a radius of curvature 1 at the origin and other points closest to the symmetry axis, with a negative sign in the eigenvalue indicating that this curvature is in an opposite sense as that described by λ_1. At the top, the radius of curvature is 3 and in the same sense as that described by λ_1, and on the sides, λ_2 vanishes because the direction of curvature is orthogonal to the tangent vector.

Now we compute the curvature integral. With ϕ the angle around the symmetry axis, the curvature integral is

$$\int_{T^2} \lambda_1\lambda_2\, dA = \int_0^{2\pi} d\theta \int_0^{2\pi} (2 - \cos\theta)\, d\phi\, \lambda_1\lambda_2 = \int_0^{2\pi} d\theta \int_0^{2\pi} d\phi\, (-\cos\theta) = 0. \tag{7}$$

Again this zero answer is generic to any surface that can be smoothly deformed to the torus. The general result (the Gauss-Bonnet formula) of which the above are examples is

$$\int_S \lambda_1\lambda_2\, dA = 2\pi\chi = 2\pi(2 - g), \tag{8}$$

where χ is a topological invariant known as the Euler characteristic and g is the genus, essentially the number of "holes" in the surface. [1] For a compact manifold with boundaries, the Euler characteristic becomes $2 - 2g - b$, where b is the number of boundaries: one can check this by noting that by cutting a torus, one can produce two discs (by slicing a bagel) or alternately a cylinder with two boundaries (by slicing a bundt cake). We will not prove the Gauss-Bonnet formula but will encounter the Euler characteristic several times in these notes.

More generally, we will encounter several examples where a topological invariant is expressed as an integral over a local quantity with a geometric significance. We now turn to a simpler example in order to allow us to introduce some basic concepts of algebraic topology.

B. Invariant integrals along paths in two dimensions: exact forms

As our first example of a topological property, let's ask about making line integrals along paths (not path integrals in the physics sense, where the path itself is integrate over) that are nearly independent of the precise path: they will turn out to depend in some cases on topological properties (homotopy or cohomology). We will assume throughout these notes, unless otherwise specified, that all functions are smooth (i.e., $\mathbb{C}^\infty$, meaning derivatives of all orders exist).

First, suppose that we deal with paths on some open set U in the two-dimensional plane $\mathbb{R}^2$. (Open set: some neighborhood of each point in the set is also in the set.) We consider a smooth path $(u(t), v(t))$, where $0 \le t \le 1$ and the endpoints may be different. (To make these results more precise, we should provide for adding one path to another by requiring only piecewise smooth paths, and require that u and v be smooth in an open set including $t \in [0, 1]$. For additional rigor, see the first few chapters of W. Fulton, "Algebraic Topology: A First Course", Springer).

Now let $f(x, y) = (p(x, y), q(x, y))$ be a two-dimensional vector field that lets us compute line integrals of this path:

$$W = \int_0^1 dt\, p\frac{du}{dt} + q\frac{dv}{dt}dt, \tag{9}$$

where p and q are evaluated at $(x(t), y(t))$.

Mathematical note: in more fancy language, f is a differential form, a "1-form" to be precise. All that means is that f is something we can use to form integrals over paths that are linear and probe the tangent vector of the path. Another way to state this, with which you may be more familiar is that the tangent vector to a path, which we call a "vector", transforms naturally in an opposite way to the gradient of a function, which we call a "covector". To

---

[1] A good question is why we write the Euler characteristic as $2-2g$ rather than $1-g$; one way to motivate this is by considering polygonal approximations to the surface. The discrete Euler characteristic $V - E + F$, where V, E, F count vertices, edges, and faces, is equal to χ. For example, the five Platonic solids all have $V - E + F = 2$.

convince yourself that this is true, think about how both transform under a linear transformation on the underlying space. We will say a bit more about such forms in a moment.

Our first goal is to show that the following three statements are equivalent: (a) W depends only on the endpoints $(u(0), v(0))$ and $(u(1), v(1))$; (b) $W = 0$ for any closed path; (c) f is the gradient of a function g: $(p, q) = (\partial_x g, \partial_y g)$; The formal language used for (c) is that f is an *exact form*: $f = dg$ is the differential of a 0-form (a smooth function) g.

Note that (c) obviously implies (a) and (b), since then $W = g(u(1), v(1)) - g(u(0), v(0))$. To show that (b) implies (a), suppose (b) is true and (a) is not. Then there are two paths γ_1, γ_2 that have different integrals but the same endpoints. Form a new path γ so that, as t goes from 0 to $\frac{1}{2}$, γ_1 is traced, and then as t goes from $\frac{1}{2}$ to 1, γ_2 is traced opposite its original direction (now you can see why piecewise smooth paths are needed if one wants to be rigorous). Then this integral is nonzero, which contradicts (b).

It remains to show that (a) implies (c). Define $g(x, y)$ as equal to 0 at $(0, 0)$, or some other reference point in U if U does not include the origin. Everywhere else, set g equal to the W obtained by integrating over an arbitrary path from $(0, 0)$ to the final point, which by (a) is path-independent. (If U is not connected, then carry out this process on each connected component.) We will show that $\partial_x g = p$, and the same logic then implies $\partial_y g = q$. We need to compute

$$\partial_x g = \lim_{\Delta x \to 0} \frac{g(x + \Delta x, y) - g(x, y)}{\Delta x}. \tag{10}$$

We can obtain g by any path we like, so let's take an arbitrary path to define $g(x, y)$, then add a short horizontal segment to that path to define the path for $g(x + \Delta x, y)$. The value of the integral along this extra horizontal segment converges to $p(x, y)(\Delta x)$, as needed.

It turns out that the above case is simple because the plane we started with is "topologically trivial." Before proceeding to look at a nontrivial example, let us state one requirement on f that is satisfied whenever f is exact ($f = dg$). The fact that partial derivatives commute means that, with $f = dg = (p, q)$, $\partial_y p = \partial_x q$. We can come up with an elegant notation for this property by expanding our knowledge of differential forms.

Before, we obtained a 1-form f as the differential of a scalar g by defining

$$f = dg = \partial_x g \, dx + \partial_y g \, dy. \tag{11}$$

Note that we now include the differential elements dx, dy in the definition of f, and that 1-forms form a real vector space (spanned by dx, dy): we can add them and multiply them by scalars. To obtain a 2-form as the differential of a 1-form, we repeat the process: writing $f = f_i dx_i$ (with $x_1 = x, x_2 = y, f_1 = p, f_2 = q$)

$$df = \sum_j \frac{\partial f_i}{\partial x_j} \, dx_j \wedge dx_i. \tag{12}$$

where the $\wedge$ product between differential forms satisfies the rule $dx_i \wedge dx_j = -dx_j \wedge dx_i$, which implies that if any coordinate appears twice, then we get zero: $dx \wedge dx = 0$. For some intuition about why this anticommutation property is important, note that in our 2D example,

$$df = (\partial_x f_y - \partial_y f_x) dx \wedge dy, \tag{13}$$

so that the function appearing in df is just the curl of the 2D vector field represented by f. So our statement about partial derivatives commuting is just the statement that if $f = dg$, then $df = 0$, or that the curl of a gradient is zero. We label any 1-form satisfying $df = 0$ a *closed form*. While every exact form is also closed, we will see that not every closed form is exact, with profound consequences.

C. Topologically invariant integrals along paths: closed forms

As an example of nontrivial topology, we would now like to come up with an example where integrals over paths are only path-independent in a limited "topological" sense: the integral is the same for any two paths that are *homotopic*, one of the fundamental concepts of topology (to be defined in a moment). Basically, two paths are homotopic if one can be smoothly deformed into another. Consider the vector field

$$f = (p, q) = \left(-\frac{y}{x^2 + y^2}, \frac{x}{x^2 + y^2} \right) = \frac{-y dx + x dy}{x^2 + y^2}, \tag{14}$$

where in the second step we have written it using our 1-form notation. This vector field is well-defined everywhere except the origin. This 1-form looks locally like the differential of $g = \tan^{-1}(y/x)$ (which just measures the angle in polar coordinates), but that function can only be defined smoothly on some open sets. For example, in a disc around the origin, the 2π ambiguity of the inverse tangent prevents defining g globally.

So if we have a path that lies entirely in a region where g can be defined, then the integral of this 1-form over the path will give the change in angle between the starting point and end point $g(u(1), v(1)) - g(u(0), v(0))$. What about other types of paths, for example, paths in $\mathbb{R}^2/\{0, 0\}$, the 2D plane with the origin omitted, that circle the origin and return to the starting point? We can still integrate using the 1-form f, even if it is not the gradient of a scalar function g, and will obtain the value $2\pi n$, where n is the "winding number": a signed integer that describes how many times the closed path $(u(t), v(t))$ circled the origin as t went from 0 to 1.

Now this winding number does not change as we make a small change in the closed path, as long as the path remains in $\mathbb{R}^2/\{0, 0\}$. What mathematical property of f guarantees this? Above we saw that any exact 1-form (the differential of a scalar function) is also closed. While f is not exact, we can see that it is closed:

$$df = \left(\partial_x \frac{x}{x^2 + y^2}\right) dx \wedge dy + \left(\partial_y \frac{-y}{x^2 + y^2}\right) dy \wedge dx = \frac{2 - 2}{x^2 + y^2} dx \wedge dy = 0. \tag{15}$$

In other words, $(-y, x)/(x^2 + y^2)$ is curl-free ("irrotational"), while $(-y, x)$ has constant nonzero curl. Now suppose that we are given two paths γ_1 and γ_2 that differ by going in different ways around some small patch dA in which the 1-form remains defined. The difference in the integral of f over these two paths is then the integral of df over the enclosed surface by Stokes's theorem, which is zero if f is a closed form.

So we conclude that if f is a closed form, then the path integral is path-independent if we move the path through a region where f is always defined. For an exact form, the integral is completely path-independent. In the case of $\mathcal{R}/\{0, 0\}$, the 1-form in Eq. 14 is locally but not completely path-independent. Both closed forms and exact forms are vector spaces (we can add and multiply by scalars), and typically infinite-dimensional, but their quotient as vector spaces is typically finite-dimensional. (The quotient of a vector space A by a vector space B is the vector space that identifies any two elements of A that differ only by an element of B). A basic object in "cohomology" is the first de Rham cohomology group (a vector space is by definition a group under addition),

$$H^1(M) = \frac{\text{closed 1-forms on } M}{\text{exact 1-forms on } M} = \frac{Z^1(M)}{B^1(M)}. \tag{16}$$

If you wonder why the prefix "co-" appears in "cohomology", there is a dual theory of linear combinations of curves, etc., called homology, in which the differential operator in de Rham cohomology is replaced by the boundary operator. However, while arguably more basic mathematically, homology seems to crop up less frequently in physics. An even simpler object is the zeroth de Rham cohomology group. To understand this, realize that a closed 0-form is one whose gradient is zero, i.e., one that is constant on each connected component of U. There are no (-1)-forms and hence no exact 0-forms. So the zeroth group is just $\mathbb{R}^n$, where n is the number of connected components.

We can show that $H^1 = \mathbb{R}$ for the unit circle S^1 using the angle form f in Eq. 14, by showing that this form (more precisely, its equivalence class up to exact forms) provides a basis for H^1. Given some other form f', we use the unit circle path, parametrized by an angle θ going from zero to 2π, to define

$$c = \frac{\int_0^{2\pi} f'}{\int_0^{2\pi} f}. \tag{17}$$

Now $f' - cf$ integrates to zero. We can define a function g via

$$g(\theta) = \int_0^{\theta} (f' - cf). \tag{18}$$

Now g is well-defined and periodic because of how we defined c, and $f' = cf + dg$, which means that f' and cf are in the same equivalence class as dg is an exact form. We say that f' and f are cohomologous because they differ by an exact form. So cf, $c \in \mathbb{R}$, generates H^1, and $H^1(S^1)$ is isomorphic to $\mathbb{R}$. With a little more work, one can show that $\mathcal{R}/\{0, 0\}$ also has $H^1 = \mathbb{R}$.

Actually we can connect the results of this section to the previous one: a general expression for the Euler characteristic is

$$\chi(M) = \sum_i (-1)^i \dim H^i(M) = \sum_i (-1)^i \dim \frac{Z^i(M)}{B_i(M)}. \tag{19}$$

The dimension of the ith cohomology group is called the ith Betti number (to be pedantic, the Betti numbers are defined for homology rather than cohomology, but we can use a duality relationship). There is a compact way to express the idea of cohomology and homology that will let us introduce some notation and terminology that comes in later. If Ω_r is the vector space of r-forms, and C_r is the dual space of r-chains, then the action of the boundary operator and the differential is as follows:

$$\begin{array}{ccccccccc} \longleftarrow & C_r & \xleftarrow[\partial_{r+1}]{} & C_{r+1} & \xleftarrow[\partial_{r+2}]{} & C_{r+2} & \longleftarrow \\ \longrightarrow & \Omega_r & \xrightarrow[d_{r+1}]{} & \Omega_{r+1} & \xrightarrow[d_{r+2}]{} & \Omega_{r+2} & \longrightarrow . \end{array} \tag{20}$$

The rth cohomology group is the quotient $\ker d_{r+1}/\mathrm{im}\, d_r$, and the rth homology group is $\ker \partial_r/\mathrm{im}\, \partial_{r+1}$.

The duality relationship is provided by Stokes's theorem. Recall that this theorem relates the integral of a form over a boundary to the integral of the differential of the form over the interior. In terms of the linear operator (f, c) that evaluates the form f on the chain c, we have the compact expression

$$(f, \partial c) = (df, c). \tag{21}$$

Now we move on to a different type of topology that is perhaps more intuitive and will be useful for our first physics challenge: how to classify defects in ordered systems.

D. Homotopy

What if we did not want to deal with smooth functions and calculus? An even more basic type of topology is homotopy theory, which can be defined without reference to calculus, differential forms, etc. (although in physics the assumption of differentiability is usually applicable). Suppose that we are given a continuous map from $[0, 1]$ to a manifold M such that 0 and 1 get mapped to the same point; we can think of this as a closed curve on M. We say that two such curves γ_1, γ_2 are homotopic if there is a continuous function (a homotopy) f from $[0, 1] \times [0, 1]$ to M that satisfies

$$f(x, 0) = \gamma_1(x), \quad f(x, 1) = \gamma_2(x). \tag{22}$$

Intuitively, f describes how to smoothly distort γ_1 to γ_2. Now homotopy is an equivalence relation and hence defines equivalence classes: $[\gamma_1]$ is the set of all paths homotopic to γ_1. Furthermore, concatenation of paths (i.e., tracing one after the other) defines a natural group structure on these equivalence classes: the inverse of any path can be obtained by tracing it in the opposite direction. (To be precise, one should define homotopy with reference to a particular point where paths start and end; for a symmetric space where all points are basically equivalent, this is unnecessary.) We conclude that the equivalence classes of closed paths form a group $\pi_1(M)$, called the fundamental group or first homotopy group. Higher homotopy groups $\pi_n(M)$ are obtained by considering mappings from the n-sphere S^n to M in the same way.

The homotopy groups of a manifold are not totally independent of the cohomology groups: for example, if $\pi_1(M)$ is trivial, then so is the first de Rham group. The examples above of $\mathbb{R}^2/\{0, 0\}$ and S^1 both have $\pi_1(M) = \mathbb{Z}$: there is an integer-valued winding number that we can use to classify paths, and this winding number can be computed by the angle form given above. So our two-dimensional examples already contains the two types of topology that occur most frequently in physics: de Rham cohomology and homotopy. We will return to homotopy in much more detail in a moment, when we explain how it can be used to classify topological defects such as vortices in broken-symmetry phases.

II. PHYSICAL PRELIMINARIES

A. Landau theory of broken-symmetry phases

Most phases of matter, including solids, magnets, superconductors, superfluids, and many others, can be understood in terms of "broken symmetry". At high temperature, fluctuations are induced by the requirement to maximize entropy, and these fluctuations tend to destroy order. As temperature is lowered, the energy gain from developing order can overwhelm the entropy gain from disorder. A remarkable fact that only became clear after the solution of the two-dimensional Ising model by Onsager in 1948 is that this energy-entropy competition can lead to a sharp phase transition, described mathematically by a singularity in some derivative of the free energy that emerges in the thermodynamic limit (the limit of an infinite number of degrees of freedom).

For our purposes, we need a way to describe such a breaking of symmetry mathematically. Rather than try to describe every microscopic degree of freedom in a complicated interacting system, we will eventually follow Landau and introduce a classical field theory in terms of some emergent or coarse-grained field that describes the type of order we wish to study. To start, let us first consider an Ising model on a hypercubic lattice. Our microscopic description is in terms of a discrete spin variable $s_i = \pm 1$ at each vertex, with the energy function

$$E = -J \sum_{\langle ij \rangle} s_i s_j. \tag{23}$$

Here ij are nearest-neighbor sites and J is some interaction strength with units of energy. At temperature $T = \infty$, the system is equally likely to be in any microstate. At $T = 0$, only two microstates occur: one with all spins up and one with all spins down. The surprise is that, if the lattice of spins is in more than one dimension, there is a nonzero temperature T_c, proportional to J, below which the zero-temperature description is qualitatively, but not quantitatively correct. As an explicit example, one can construct the lattice mean-field description of the Ising model, leading to the equation

$$m = \tanh(zJm/(k_B T) + h) = \tanh(\beta z J + h), \tag{24}$$

where $m = \langle s_i \rangle$ is the average spin, $\beta = 1/k_B T$, z is the number of nearest neighbors, h is a possible external magnetic field (defined to include a factor β), and k_B is Boltzmann's constant; for a derivation, cf. Chaikin and Lubensky. The behavior of this self-consistent equation changes at $k_B T = zJ$, and the mean-field transition temperature is $T_c = zJ/k_B$ or $\beta_c = (zJ)^{-1}$.

Physicists say that the system "breaks symmetry" below T_c and picks out a particular sign of the average spin m, where the angle brackets denote thermal averaging. Mathematicians have more satisfactory definitions (because m strictly speaking is always zero): one can look at either whether there is a nonzero correlation function $\langle s_i s_j \rangle$ as i and j become infinitely far apart, or look for a singularity in some derivative of the free energy (in a first derivative for a first-order transition, in some higher derivative for a second-order transition). Even if m is always zero in terms of the Boltzmann sum, physical systems do actually break symmetry, chiefly for dynamical reasons: for example, a bar magnet of iron will in principle explore the whole phase space and flip its north and south poles, but the time it takes to do so may be larger than the age of the universe. Hence we will mostly be content to discuss broken symmetry as real, e.g., $m \neq 0$ in the Ising model, even if that is somewhat sloppy mathematically.

A powerful way to understand the broken symmetry is in terms of two symmetry groups: G, the high-temperature symmetry group, and H, the residual symmetry group that survives in the low temperature phase. We define the "order parameter manifold" as the quotient $M = G/H$, where "dividing by H" (taking cosets of H in G) means that we identify two elements of G that differ by an element of H; note that this is not in general a group. The notion of an "order parameter" is basic in Landau theory: it is the field we use to model all the complicated microscopic states in terms of one, or a few, macroscopic variables. The idea of the order parameter manifold is that, for many interesting phenomena, we do not care especially about the magnitude of the order parameter itself. We care instead about the set of distinguishable low-temperature states at an arbitrary temperature in the ordered phase, which is exactly M.

One reason the set M is important will become clear when we discuss topological defects in the next chapter. Topological defects in an ordered phase can be classified using mappings from spheres in real space to the order parameter manifold M, i.e., the homotopy groups $\pi_n(M)$. We will explain this result and see a number of examples; two other reasons the manifold M is important in practice is that, since moving from one point in M to another is naïvely a "soft" or massless fluctuation, while changing the magnitude of the order parameter is a "hard" or massive fluctuation, using a field-theory description that involves only the degrees of freedom in M, known for historical reasons as a "nonlinear σ-model", is frequently useful.

Landau theory, which we will continue to discuss in the following section, is not restricted to second-order transitions like those we have described above. Actually one of the most powerful predictions of Landau theory, questioned theoretically in recent years, is that second-order transitions require a change of symmetry: more specifically, the symmetry group H must be a proper subgroup of G or vice versa. Otherwise, Landau theory predicts a first-order phase transition (first-order transitions are always allowed by symmetry, and can be understood simply as level crossings in the free energy of two different phases).

B. Quantum and statistical path integrals

There are two field theories we will deal with in our treatment of broken-symmetry phases. The first is Landau-Ginzburg theory, which can be understood as a power series and gradient expansion of the energy density in terms of

an order parameter field. For the example of the Ising model, in zero magnetic field we can expand the energy density in even powers of $m(\mathbf{r})$:

$$Z_{\mathrm{LG}} = \int (Dm)\, e^{-\beta \int d^d\mathbf{r}\, a_0 + a_1 m^2 + a_2 m^4 + \ldots + b_0 (\nabla m)^2 + \cdots}. \tag{25}$$

Here the measure of the integral can be defined more precisely in Fourier space, where omitting high-wavevector components is typically necessary for a sensible theory (in condensed matter, this makes physical sense as a "short-distance cutoff" below which the field m is not meaningful). This is a vastly simplified version of the original problem in at least two ways: we are assuming that the integration over the coarse-grained order parameter field $m(\mathbf{r})$, and we are not doing any microscopic calculation of the coefficients that appear in the expansion.

A remarkable fact is that the above Landau-Ginzburg theory can be not just qualitatively correct but actually exact for some properties, such as "critical exponents" near second-order phase transitions, even without a microscopic calculation of the coefficients. Such properties are referred to as "universal": universal properties depend on symmetry and dimensionality but little else. For example, the liquid-gas critical point in the phase diagram of water has the same critical exponents as the Ising phase transition in three dimensions. As a mathematical example of where universality comes from, the terms beyond a_2 in the above energy turn out not to impact critical exponents and selected other properties, as long as the lower-order coefficients have appropriate signs. We will study one or two examples of critical points later, concentrating on examples where topological considerations are important.

The Landau-Ginzburg theory could have been motivated as a high-temperature expansion: we are justified in concentrating on low powers of m because at high enough temperature the order parameter should not be too large. An alternate field theory, the nonlinear σ-model, can be viewed as an expansion starting from zero temperature. We will concentrate on systems in which the zero-temperature phase breaks a continuous symmetry, so that M is a continuous manifold; this includes, for example, Heisenberg and XY magnets, in which $M = S^2$ and $M = S^1$ respectively, but not Ising magnets, where M is a set with two elements. For an XY magnet [2], we can label a ground state simply by an angle θ between 0 and 2π (the order parameter has a magnitude Δ as well, but all ground states have the same magnitude of the order parameter by symmetry).

When the temperature is slightly increased, fluctuations of the order parameter will take place. The nonlinear σ-model is a theory that ignores fluctuations of the order parameter *magnitude* but captures fluctuations in its *direction*, which are lower in energy or "softer". More precisely, the nonlinear σ-model into a symmetric space $M = G/H$ is defined as a path integral over an M-valued field. For the XY case above, this can be written simply in terms of a spatially varying angle $\theta(\mathbf{r})$:

$$Z_{\mathrm{NL}\sigma\mathrm{M}} = \int (D\theta)\, e^{-\int d^d\mathbf{r}\, g(\nabla\theta)^2}, \tag{26}$$

where we have incorporated β into the definition of the coupling g. We will return to this model once we have said a bit more about topological defects; it turns out that for our XY example in two spatial dimensions, the physics depends crucially on "vortices", and in fact shows a phase transition that would not be present if, hypothetically, the field θ were not periodic and Eq. 26 became just the Gaussian model.

We have written both of the above field theories in a classical or Euclidean representation, where A natural question is how the partition function integral in such a theory is related to the quantum path integral that may be familiar from an advanced course in quantum mechanics. The easiest example of the analytic continuation to imaginary time that connects the two types of path integrals is for the harmonic oscillator. Its partition function at a finite temperature T is

$$Z_{\mathrm{harmonic}} \approx \int dx(\tau)\, e^{\int_0^\beta d\tau\, \dot{x}^2(\tau)/2m + kx^2/2}, \tag{27}$$

where there are periodic boundary conditions on $x(\tau)$: $x(\beta) = x(0)$. A worthwhile calculation (hint: simplify the integral by considering Fourier components of $x(\tau)$) leads to the result

$$Z_{\mathrm{harmonic}} = \frac{1}{2\sinh(\beta\hbar\omega/2)} = \sum_{n=0}^{\infty} e^{-\beta\hbar\omega(n+1/2)}, \tag{28}$$

---

[2] We choose the example of $M = U(1) \cong SO(2)$ here for a reason. It turns out that, for the nonlinear σ-model to include gapless excitations, the form of the theory becomes more complicated. For Lie groups more complicated than $U(1)$, an additional term of topological origin is required, leading to the Wess-Zumino-Novikov Witten model that we discuss in Section .

where the last expression is what we would calculate from the spectrum. Now analytically continuing this calculation from imaginary time τ gives a trace of the

$$Z_{\text{harmonic}} = \text{Tr } e^{-\beta H} \to \text{Tr } e^{-itH/\hbar} = \int dx_0 \, U(x_0, t; x_0, 0), \tag{29}$$

where in the last step we have used the position basis to put the result in terms of matrix elements of the unitary time evolution operator U. Now the divergence of Z at real times $t = 2\pi n/\omega$, for integer n, can be simply interpreted: at these times all the energy eigenstates that appear in an arbitrary initial condition appear with exactly the same phases, so the state is (aside from an overall phase factor) exactly the initial state, the time evolution operator is the identity, and the trace diverges. We now turn to a beautiful geometric property of quantum mechanics that was understood relatively recently: the geometric or Berry phase.

C. Berry phases in quantum mechanics

An important result from undergraduate quantum mechanics is the "adiabatic approximation". Suppose that a system is prepared in a nondegenerate eigenstate of a time-dependent Hamiltonian H. For later reference, we will write H as a function of some parameters λ_i that depend on time: $H(t) = H(\lambda_1(t), \lambda_2(t), \ldots)$. If the eigenstate remains nondegenerate, then the adiabatic approximation is the result that if the Hamiltonian changes slowly in time (how slowly depends primarily on the energy gap between adjacent eigenstates), then there are no transitions between eigenstates.

This approximation, when correct, actually only gives part of the story: it describes the probability to remain in the eigenstate that evolved from the initial eigenstate, but there is actually nontrivial information in the *phase* of the final state as well. This result may seem quite surprising because the overall phase in quantum mechanics is in general independent of observable quantitites. However, the Berry phase from an adiabatic evolution is observable: for example, one system can be taken around a closed path in parameter space, while another system initially identical to the first can be taken around a different path, or the null path; an interference experiment on the final states will reveal the Berry phase. The first example of this type of geometric phase in physics was found more than fifty years ago by Pancharatnam in an optical example, but the classic Berry paper of 1984 was the first to explain the concept in its full generality.

Berry's result for a closed path is relatively simple to state, but some careful thought is required to understand and derive it. In moving a system adiabatically around a closed path in parameter space, the final wavefunction is in the same eigenstate as the initial one (again, under the assumptions of the adiabatic approximation as stated above), but its phase has changed:

$$|\psi(t_f)\rangle = e^{-(i/\hbar)\int_{t_i}^{t_f} E(t')\,dt'} e^{i\gamma} |\psi(t_i)\rangle. \tag{30}$$

Here $E(t')$ means the corresponding eigenvalue of the Hamiltonian at that time, and γ is the Berry phase, expressed as an integral over a path in *parameter* space with no time-dependence:

$$\gamma = i \int \langle \psi(\lambda) | \nabla_\lambda | \psi(\lambda) \rangle \cdot d\lambda. \tag{31}$$

Note that there are two different arguments of ψ in the above formulas. When ψ has a time argument, it means the wavefunction of the system at that time. When ψ has a parameter argument, it means the wavefunction we have chosen of the Hamiltonian . A key assumption of the Berry phase derivation is that there is some smooth choice of the $\psi(\lambda_i)$ throughout a surface in parameter space with the loop as boundary.

For an open path, we need to describe the phase of the wavefunction relative to this reference set, so the expression becomes more complicated (for the closed path, we could simply compare the initial and final wavefunctions, without needing the reference set at these points). We will show that, assuming $\psi(t_i) = \psi(\lambda(t_i))$ so that the initial wavefunction is equal to the reference state at the corresponding value of parameters,

$$\langle \psi(\lambda_i(t)) | \psi(t) \rangle = e^{-(i/\hbar)\int_0^t E(t')\,dt'} e^{i\gamma}, \tag{32}$$

i.e., the Berry phase appears when comparing the actual time-dependent evolved state $\psi(t)$ to the reference state at the same point. We can take the time derivative of both sides and use the time-dependent Schrödinger equation

$$i\hbar \frac{\partial \psi}{\partial t} = H(t)\psi. \tag{33}$$

The two sides agree initially if we choose the proper boundary condition on the Berry The time derivative of the left side is

$$\langle\psi(\lambda_i(t))|\frac{-iE(t)}{\hbar}|\psi(t)\rangle + \frac{d\lambda_j}{dt}\langle\partial_{\lambda_j}\psi(\lambda_i(t))|\psi(t)\rangle, \tag{34}$$

so writing $e^{i\theta(t)} = \langle\psi(\lambda_i(t))|\psi(t)\rangle$, we have computed

$$\frac{d}{dt}e^{i\theta(t)} = \left(\frac{-iE(t)}{\hbar} + \frac{d\lambda_j}{dt}\langle\partial_{\lambda_j}\psi(\lambda_i(t))|\psi(\lambda_i(t))\rangle\right)e^{i\theta(t)}, \tag{35}$$

and this is satisfied if (note that for E we do not need to distinguish between time and λ dependent)

$$\dot{\theta}(t) = -\frac{E(t)}{\hbar} - i\frac{d\lambda_j}{dt}\langle\partial_{\lambda_j}\psi(\lambda_i(t))|\psi(\lambda_i(t))\rangle, \tag{36}$$

which is our desired conclusion after noting that (dropping the explicit t dependence)

$$\frac{d\gamma}{dt} = i\frac{d\lambda_j}{dt}\langle\psi(\lambda_i)|\partial_{\lambda_j}\psi(\lambda_i)\rangle = -i\frac{d\lambda_j}{dt}\langle\partial_{\lambda_j}\psi(\lambda_i)|\psi(\lambda_i)\rangle. \tag{37}$$

The negative sign in the last equality, which shows that the "Berry connection" or "Berry vector potential" $A_j = i\langle\psi(\lambda_i)|\partial_{\lambda_j}\psi(\lambda_i)\rangle$ is real, follows from noting that $\partial_{\lambda_j}\langle\psi|\psi\rangle = 0$ by constancy of normalization. Crucial in this derivation was the evolution of H, which means that the evolution is more than just a simple phase factor, even though the actual rate of change in H drops out and only the path taken by H enters the Berry phase.

Now one can ask whether the Berry phase is independent of how we chose the reference wavefunctions (in this case, the $U(1)$ degree of freedom in the wavefunction at each λ). While for an open path it clearly is not phase-independent, the Berry phase is phase-independent for a closed path, for exactly the same reasons as a closed line integral of $\mathbf{A}$ is gauge-independent in electrodynamics: we can integrate the "Berry flux" or "Berry curvature" $\epsilon_{ij}\partial_i A_j$ (which you can check is phase-independent, just like $F_{\mu\nu}$ in electrodynamics) on the surface bounded by the path. Alternately, we can note that a phase change changes A by the gradient of a scalar, so that on a closed loop, there is no change.

To get some geometric intuition for what the Berry phase means, we explain why the Berry connection A is sometimes called a connection, and the flux F is sometimes called a curvature. A connection is a way to compare vector spaces that are attached to different points of a manifold, forming a "vector bundle". In our case, there is a one-dimensional complex vector space attached at each point in parameter space, spanned by the local eigenstate. The inner product lets us compare vectors at the same point in parameter space, but the Berry connection appears when we try to compare two vectors from slightly different points.

A useful example of a real vector bundle is the "tangent bundle" to a Riemannian manifold (say, a sphere), made up of tangent vectors at each point, which have a dot product corresponding to the inner product in quantum mechanics. The connection in this case, which gives rise to "parallel transport" of tangent vectors, is related to the same curvature that we previously discussed with the Gauss-Bonnet theorem. Consider an airplane moving around the surface of the Earth and carrying a gyroscope that is fixed to lie in the tangent plane to the Earth's surface (i.e., free to rotate around the normal axis to the tangent plane). If the airplane follows a great circle, then it will appear to be going straight ahead to a passenger on board, and the gyroscope will not rotate relative to the plane's axis.

However, if the airplane follows a line of latitude other than the equator, or any other path that is not a "geodesic" (see a differential geometry text for details), it will feel constantly as though it is turning, and the gyroscope will appear to rotate relative to the airplane's direction. After going around a closed path in the airplane, the gyroscope may have rotated compared to a stationary gyroscope (the same physics that underlies Foucault's pendulum). As an exercise, you can work out that the total angle of rotation in circling a line of latitude is $2\pi\sin(\phi)$, where ϕ is the latitude. At the equator this gives no rotation, while at the north pole this gives a 2π rotation. This is a geometrical version of the same idea of holonomy (failure of a gyroscope to return to its initial direction) that underlies the Berry phase.

Note that a vector potential in a gauge theory and the associated Wilson loop are also examples of the concept of holonomy in a (now complex) vector bundle. The $U(1)$ Berry phase described above generalizes immediately to a non-Abelian Berry phase, which has some important applications that were only recently discovered.

III. TOPOLOGICAL DEFECTS

(For this week of the course, I will follow the review article by Mermin, available on the course internal website, and hence will not prepare lecture notes except for topics that are not in Mermin. The only material on individual

topological defects we covered that is not in Mermin was an example: spinor Bose condensates for $s = 1$ have "half-vortex" topological defects, even though $\pi_1(M) = \mathbb{Z}$ just as in an ordinary superfluid. These notes focus on collective physics of topological defects.)

A. Collective properties of topological defects I: the Kosterlitz-Thouless transition

One of the most remarkable examples of a collective effect arising from many topological defects is the superfluid transition in two spatial dimensions. We sketch a theoretical prediction by Kosterlitz and Thouless, anticipated in part in previous work of Berezinskii, that received spectacular experimental confirmation in work of Bishop and Reppy on ^{4}He films. Our starting point is the two-dimensional "XY model": the local spin variable on each site of a lattice is a unit vector on the circle, with the lattice Hamiltonian

$$H = -J \sum_{\langle ij \rangle} \mathbf{s}_i \cdot \mathbf{s}_j = -J \sum_{\langle ij \rangle} \cos(\theta_i - \theta_j), \tag{38}$$

where J is an energy and the sum is over nearest-neighbor pairs. In the second equality we have introduced an angle θ via $s_x + i s_y = e^{i\theta}$.

This model has the same symmetry as the superfluid transition in a film of atoms with no low-temperature internal degrees of freedom, by the following argument: the Bose condensation transition means that one quantum state has a macroscopic number of atoms, and the wavefunction of this state $\psi(\mathbf{r})$ can be taken as the order parameter in the ordered phase. Suppose that we are at low temperature so that the spin moves only slightly from one site to the next. Then, in going around a large circle, we can ask how many times the spin winds around the unit circle, and define this as the "winding number" $n \in \mathcal{Z}$. Note that if the winding number is nonzero, then the continuum limit must break down at some point within the circle, as otherwise we would have the same angular rotation $2\pi n$ around circles of smaller and smaller radius, implying larger and larger gradients and hence infinite energy (we will calculate the energy of a vortex below).

This will also let us see from a fairly simple calculation how there can be continuously varying exponents in the power-law correlations of the 2D XY model at low temperature. The assumption we'll need to make is that "vortices" are unimportant at sufficiently low temperature, so that the 2π periodicity of the phase can be ignored (a vortex in 2D is a point where the magnitude of the order parameter vanishes, around which the phase of the order parameter changes by a multiple of 2π). This calculation can be justified by looking at the RG flow in a space of two parameters: the temperature, and the vortex fugacity (essentially a parameter controlling how many vortices there are). An excellent reference for this RG flow is the original paper by J. M. Kosterlitz and D. J. Thouless, J. Phys. C **6**, 1181 (1973). If there are no vortices and the magnitude of the order parameter is constant, then the effective partition function is, where θ is no longer restricted to be periodic,

$$Z = \int D\theta(r)\, e^{-\frac{K}{2}\int (\nabla\theta)^2\, d^2 r} \tag{39}$$

Here K is a dimensionless coupling incorporating temperature that, in a lattice model such as (38), can be obtained by linearization. We can define a "superfluid stiffness" with units of energy $\rho_s = (k_B T)K$ that measures the energy required to create a twist in the superfluid phase. One way to look at this nonlinear sigma model is as describing slow variations of the ordered configuration at low temperature: the magnitude is fixed because fluctuations in magnitude are energetically expensive, but since there are degenerate states with the same magnitude but different θ, slow variation of θ costs little energy.

In the model with no vortices, i.e., with θ treated as a real-valued rather than periodic field, the spin correlation function, which we will find goes as a power-law, is

$$\langle \mathbf{s}(0) \cdot \mathbf{s}(r) \rangle = \mathrm{Re} \langle e^{i\theta(0)} e^{-i\theta(r)} \rangle. \tag{40}$$

(Actually taking the real part is superfluous if we define the correlator of an odd number of θ fields to be zero.) We choose not to rescale the θ field to make K equal to unity, since such a rescaling would modify the periodicity constraint $\theta = \theta + 2\pi$ in the model once vortices are restored. To get this correlation function, we first compute the correlation of the θ fields $G(r) = \langle \theta(0)\theta(r) \rangle$, which we will need to regularize by subtracting the infinite constant $G(0) = \langle \theta(0)^2 \rangle$.

From our previous experience with Gaussian theories, we know to write $G(r)$ as an integral over Fourier components:

$$G(r) = \frac{1}{(2\pi)^2 K} \int^{a^{-1}} \frac{e^{i\mathbf{k}\cdot\mathbf{r}}}{k^2}\, d\mathbf{r}. \tag{41}$$

However, this integral is divergent at $k = 0$ ("infrared divergent"; "ultraviolet divergent" means at $k = \infty$). We can regularize the divergence by subtracting out the formally infinite quantity $G(0)$, and then using $\tilde{G}(r) = G(r) - G(0)$ to calculate physical quantities.

$$\tilde{G}(r) = -\frac{1}{(2\pi)^2 K} \int^{a^{-1}} \frac{1 - e^{i\mathbf{k}\cdot\mathbf{r}}}{k^2}\, d\mathbf{r}. \tag{42}$$

Now the integral can only be a function of the ratio r/a (you can check this by changing variables). As $a \to 0$, the leading term in the integration is proportional to

$$\tilde{G}(r) = \frac{1}{(2\pi)^2 K} \int^{a^{-1}} \frac{dk}{k} = -\frac{\log(r/a)}{2\pi K} + \ldots, \tag{43}$$

where one factor of 2π was picked up by the angular integration. Note that at large r, $\tilde{G}$ is divergent, which makes sense since θ is unbounded.

We now want to calculate the resulting spin-spin correlator. To do this we need to use a fact about Gaussian integrals. Recall the elementary Gaussian average

$$\langle e^{iJ\phi} \rangle = (A/\sqrt{2\pi}) \int_{-\infty}^{\infty} d\phi\, e^{iJ\phi} e^{-\frac{1}{2}A\phi^2} = e^{-\frac{1}{2}A^{-1}J^2} = e^{-\frac{1}{2}\langle J^2\phi^2 \rangle}. \tag{44}$$

This generalizes to the continuum limit in the following way (left to reader):

$$\langle e^{-i\theta(r)+i\theta(0)} \rangle = e^{-\frac{1}{2}\langle (\theta(r)-\theta(0))^2 \rangle} = e^{G(r)-G(0)}. \tag{45}$$

So finally

$$\langle \mathbf{s}(0) \cdot \mathbf{s}(r) \rangle = \mathrm{Re} \langle e^{i\theta(r)-i\theta(0)} \rangle = \frac{1}{(r/a)^{1/2\pi K}}. \tag{46}$$

We expect on physical grounds that this power-law correlation function ("algebraic long-range order") cannot survive up to arbitrarily high temperatures; above some maximum temperature, there should be a disordered phase with exponentially decaying correlations. To understand how our physical expectation of exponentially short correlations at high temperature is made correct, we give a simple picture due to Kosterlitz and Thouless (which is supported by a more serious RG calculation). The picture is that the phase transition results from an unbinding of logarithmically bound vortex-antivortex pairs, which can be viewed as the plasma-gas transition of a two-dimensional Coulomb plasma [3]. Vortex-antivortex pairs are logarithmically bound because the energy of a single vortex of winding number n goes as, restoring a factor of temperature in order to obtain units of energy,

$$E = \frac{1}{2} K (k_B T) \int_a^L (n/r)\, d^2 r \sim \pi n^2 K \log(L/a). \tag{47}$$

Here L is the long-distance cutoff (e.g., system size) and a is the short-distance cutoff (e.g., vortex core size). Although the energy of a single vortex in the infinite system diverges, the interaction energy of a vortex-antivortex pair does not; each vortex has energy given by (**??**), but with the system size replaced by the intervortex spacing. Note that changes of order unity in the definition of this spacing or the core size will add constants to the energy but not change the coefficient of the logarithm.

We would like to compare the free energy of two phases: one in which vortices are "bound" in pairs, and essentially do not modify the Gaussian model, and one in which vortices are numerous and essentially free, although the system is still charge-neutral (total winding number 0). Suppose the vortices in the free phase have typical separation L_0. Then each vortex can be distributed over a region of size ${L_0}^2$, and the entropic contribution to the free energy per vortex is $-TS = -2T \log L_0/a$. The energy cost is $E = \pi n^2 J \log L_0/a$, so there should be a phase transition somewhere near $T_{KT} = \pi J/2k_B$, where we have written $K = J/k_B T$ in order to define a coupling energy scale J.

---

[3] The phase with unbound vortices can be viewed as a "plasma" phase since it is has unbound positive and negative charges (vortices) as in a plasma. Note, of course, that the logarithmic interaction between vortex charges in 2D is just like that between Coulomb charges in 2D. A more serious calculation constructs an RG flow in terms of vortex fugacity to show that below a critical temperature vortices are irrelevant (their fugacity scales to 0), while above that temperature vortices are relevant (their fugacity increases upon rescaling).

A bit more work shows that this coupling scale, as we have defined it, is exactly the "superfluid stiffness" ρ_s that measures the energy induced by a twist in the superfluid phase. More precisely, the Kosterlitz-Thouless transition occurs when the asymptotic long-distance stiffness ρ_s^∞, including renormalization by bound vortex pairs, satisfies

$$\rho_s^\infty = \frac{2k_B T_{KT}}{\pi}. \tag{48}$$

This prediction of a universal jump at T_{KT} in the superfluid stiffness was beautifully confirmed in experiments by Bishop and Reppy. Another way to state this result is that a 2D superfluid that starts at short distances with stiffness less than this value allows vortices to proliferate and reduce the long-distance superfluid stiffness to zero. This behavior is rather different from that in higher dimensions, where the superfluid density flows to zero smoothly.

B. Collective properties of topological defects II: the vortex lattice

In the above example, finite temperature induces a pattern, because of entropic considerations, that on average has equal numbers of vortices with positive and negative winding numbers ("vortices" and "anti-vortices"). Rotating a three-dimensional superfluid system generates line vortices along the direction of rotation that, at sufficiently low temperature, form an ordered lattice. [4] These line vortices can be viewed as arising from the boundary conditions at infinity, which must describe a nonzero overall flow. A key property of superflows is that the flow is "locally irrotational": if the superfluid flow can be defined throughout in a closed region, then the total angular momentum is zero. However, the existence of vortex cores allows nonzero winding number: more precisely, the circulation around a vortex of winding number n is quantized to be

$$\oint \mathbf{v} \cdot d\mathbf{l} = \frac{2\pi\hbar}{m}. \tag{49}$$

The density of vortices in a rotated superfluid is determined by the condition that the overall circulation match that imposed by the rotation. [5]

This vortex lattice, first studied theoretically by Tkachenko, is a slightly simpler analogue of the Abrikosov vortex lattice generated by applying a magnetic field to type-II superconductors. The main difference is that a superconductor has two inequivalent energy scales: the correlation length ξ, which describes spatial variations of the order parameter, and the penetration depth λ, which describes spatial variations of the gauge field. In neutral superfluids the second term is absent and microscopic details only enter in one length scale, the vortex core size a (the analogue of ξ in a superconductor).

Tkachenko's approach was to solve for the full superfluid flow field using elliptic functions. Here we give a trick to compare the energy of different vortex lattices and understand why the triangular lattice is favored in the dilute limit (the distance between vortices is much larger than the core size). The third dimension can be ignored at low temperature where it is favorable for vortex lines to be straight.

(to be continued)

---

[4] Strictly speaking this lattice will, in a system of finite thickness, not have truly long-ranged correlations at nonzero temperature, but rather algebraic decay, because of the Mermin-Wagner theorem.

[5] Note that the velocity of the rotating container must exceed the first critical velocity in order for the vortices to be created, if the superfluid was initially stationary. Otherwise the superfluid can just remain stationary, as a moving wall with velocity less than the critical velocity will not excite particles out of the superfluid.